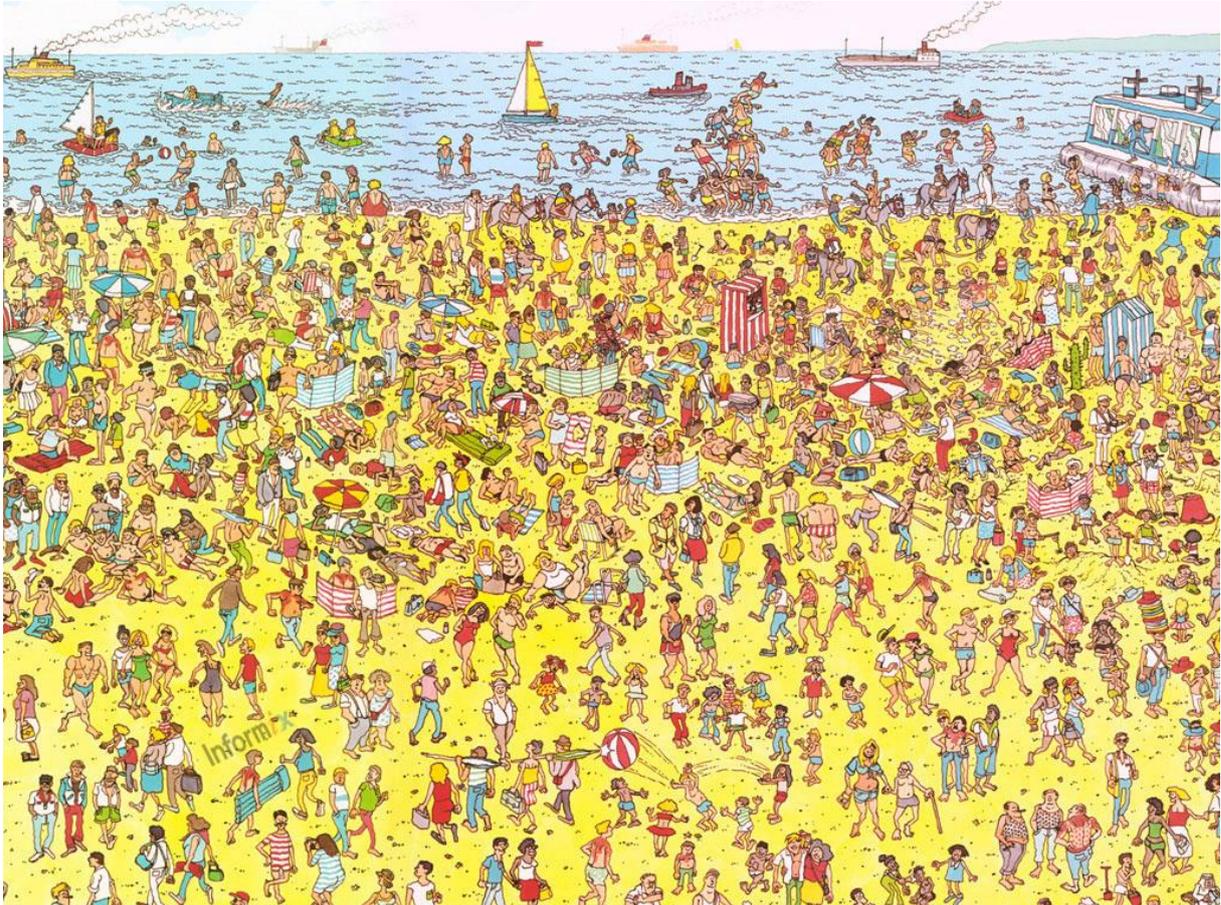


Where is Informix ?



« Tales of the best DBMS on the marketplace, that everybody is reporting as missing, except the users community and emerging industries »

Pont l'Abbé (France), March 5th 2014





How many times have I heard this sentence during the last years: "Informix? Didn't this product disappear in 2001", or even "Informix? I don't know this name. Is this a Gaulish comics character?"

I won't hide that every time I hear this, the latin part of my blood just starts boiling, while the Gaulish one calls for revolution. But while we are talking about behavioral psychology, let's stick to the Anglo-Saxon factual approach jointly with the Germanic rigor to analyze more seriously the reasons why we can hear this kind of silly statements.

The Informix Golden Era

No one will deny that in the 80's and 90's, Informix was part of the "gang of the four DBMS vendors", that is Informix, Sybase, Ingres and O...cle. At this time, the Unix operation system was taking the leading part of the market, mainly because of the drastic cost decrease as compared with the very expensive mainframe platforms. Remember that at this time, all the analysts thought and said that the mainframe universe was living its very last years (as far as I know, many of them are still up and running). Mike Saranga had shown the path to follow for the RDBMS with DB2, the same way that Roger Sippl and Laura King followed and developed brilliantly the very first versions of C-ISAM, RDS, finally Informix. Those were fat years for the IT industry in general, companies and administrations had no budget restrictions and did not hesitate in investing massively in that industry.

Those were the most beautiful years for Informix Software, which certainly had a technology leadership in the DBMS world conquered thanks to its real multi-threaded architecture (the Dynamic Scalable Architecture aka DSA), along with a head position in the marketplace. At this time, this technology leadership has been (maybe too much) ahead the market needs, as the real added value of many of the Informix new features has not been understood as it should have. We'll talk more about this topic later.

The integration of Illustra, first ODBMS (object oriented created by the visionary Michael Stonebraker, who later gave birth to PostgreSQL), combined performance, robustness and the unmatched scalability of the multi-threaded architecture with the capacity to handle all kind of heterogeneous existing or non-existing data types thru the use of Datablades like TimeSeries, Spatial, Text Search, MQ Series, Image/Video, C-ISAM, Excalibur Text Search. In case you had to create your custom data types and the matching SQL, the Datablade Developer Toolkit had all the necessary tools to create them.

The decline of Informix Software Inc.

The "stupid bizarre accounting practices" leading to a severe drop of the IFMX stock share have certainly contributed to give a negative image to Informix Software, fact that the competitors largely used to kill a solid technologic reputation built over years. Were they as white (no joke here) as they claimed, I have no clue: "no evil seen, no evil done".

The weak position of the IFMX stock allowed IBM Corporation to acquire the DBMS and language part of Informix in 2001 for the amount of 1 billion US\$.

What were the main intentions of the IBM executive management behind this acquisition? Were they to grab the technology and put it into DB2? Were they to kill a serious competitor? Were they to buy a customer portfolio and migrate it to DB2? I do not have this answer, but would guess it was somewhere in-between the three options. Something certain was that IBM decided to incorporate many parts of Informix into the DB2 code. Luck or no luck, this project of integration has been cancelled soon after it was initiated for technical reasons...

Storm warning on the Informix brand.

Surprised and probably disappointed by this unexpected project cancellation, the IBM management and sales force did not develop nor apply a clear strategy in the short, mid and long term regarding the Informix database engine continuation. This has certainly been the main argument for the competitors to stab Informix, causing severe casualties to the product fate. Customers felt kind of hijacked by the IBM sales force, without having in hands clear elements to decide about the future of their IT infrastructure.

The Informix customers also did not receive positively this erratic attitude, mainly because they did not understand why unknown people suddenly appearing in their landscape would force them to make radical changes in their DBMS strategy, although they were totally satisfied with what they had. How come this great DB engine could be living its last few days?

It was also rumored that at this time some IBM sales persons would not hesitate to sell the 'Red DBMS' to Informix customers in case they refused to go DB2, arguing that Informix was already a 'has-been' product and keeping it was too risky for their companies.

Changing DBMS brand is never an easy decision to take for a company, whatever its size is, and even worse if everything is running smoothly and no one is complaining. Wasn't the sweet mermaid song of the competitors building up the path to over budgets, penalized users, decreased productivity and finally unhappy executives due to financial losses? This is certainly the right time to ask a few questions arising from proven facts:

- Why an important share of the Informix applications having migrated to competitors have come back to Informix after a few months or a few years?
- How many opposite scenarios have occurred? (i-e competitor to Informix, then back to competitor)
- How many Informix applications have migrated to competition due to technical or financial dissatisfaction? (not considering GPU Licenses in financial factors though)
- How many Informix applications have migrated to competition based on purely executive/political decision, just because "the management has decided", ignoring financial, technical and user satisfaction criteria?

These legitimate and common sense questions finally gave birth to an insurgency movement among the Informix users community. This is the refusal message that the IIUG (International Informix User Group, or www.iiug.org) has been delivering to IBM on behalf of the community.

IBM, under tight pressure, finally had to give up the idea of killing Informix.

IBM understands the message from the customers: Informix will proceed

Literally constrained by the IIUG, IBM decided to maintain the product, which quickly, against all odds for the IBM executive management, appeared to be a profitable product line. One thing leading to another, up to seven versions of Informix have been released since the acquisition, that is:

- version 9.30 in 2001, including 7 maintenance releases,
- 9.40 in 2003, including 9 maintenance releases,
- version 10.0 in 2005, including 11 maintenance releases,
- 11.10 in 2007, including 3 maintenance releases,
- 11.50 in 2008, including 9 maintenance releases,
- 11.70 in 2010, reaching 8 maintenance releases still ongoing,
- and finally 12.10 in 2013, actually at its third maintenance release.

This makes a total of 48 product releases in 13 years, which is more characteristic of a healthy product than a dead or discontinued product. Let us check this history more in detail.

Informix 9.30: completion of the Illustra code merge.

Starting with « OK, if IIUG stops us from discontinuing Informix, let's keep the focus on maintenance and see what happens » in mind during the period covering IFMX 9.30 and IFMX 9.40, things started moving with version 10.0 which contained significant enhancements, like index self-join, configurable page size, and an important one being breaking the barrier of the 2Gb maximum size for a chunk. Enterprise Replication administration tools have been extended. We can also not a number new of SQL statements or syntax, initiating a long synchronization process for the SQL syntax with the competitors. Nevertheless, these new features are “shy” and do not demonstrate heavy investments from IBM management.

Informix 11.10 : big investment in R&D

A consequent involvement is notable in the 11.10 version: this version is much richer in terms of significant innovations like the introduction of the MACH11 cluster, featuring the Remote Standalone Secondary (RSS) and the Shared Disk Secondary (SDS). This concept open the way to a real cluster architecture which nowadays not only has nothing to envy of competitors, but certainly the competitors would love to have. Important performance enhancements are also implemented, as the eviction of blocking checkpoints, online index creation, direct-IO mode for cooked-file chunks and more. This version finally introduces Open Admin Tool, the GUI Informix administration tool, the capacity for SQL statements performance deep analysis at runtime with SQLTRACE. Also an internal Informix scheduler has been deployed, whose purpose is to handle admin and monitoring tasks in an unattended manner.

Informix 11.50 : small enhancements and adjustments, no revolution.

Looking at Informix 11.50 will bring the conclusion that it is more a version for stabilization version, only introducing noiseless but efficient enhancements. Data compression for tables is part of those enhancements, although it brings a clear advantage when dealing with data storage budgets, with an IO performance increasing as a side-effect. As for the rest of the new features, one cannot really feel a major investment from the database vendor.

Informix 11.70 : IBM believes and invests in Informix

The first release of 11.70 counts no less than 78 new features and enhancements. The most significant one is probably the introduction of the Flexible Grid. Based on the MACH11

architecture, the Flexible Grid is a new and non-complex way of creating and administrate a cluster of Informix DB servers. Beyond its child-like simplicity, the Grid allows operations undone with Informix (and many others in this price range) until this time, such as connecting heterogeneous OS and/or hardware based Informix Servers, mixing in a very flexible way Enterprise Replication, HDR, RSS and SDS, all into one "Grid". With the Informix Flexible Grid, it is henceforth possible to perform a version upgrade in a production environment with ZERO DOWNTIME, which is one of the many tricky operations that you can execute with this non-complex to administrate Database Servers Cluster.

Other new features like the Automatic Storage Pool will confirm the direction taken by IBM towards the auto-administration of IDS, which is completed by a number of auto-tunable engine parameters coming out at each new release. No doubt that these features have been determining factors of the decision taken by leaders of the Network Industry to choose Informix to embed in their new telephone switchboards, and also by house automation industry, and now mobile devices industry made possible because Informix works on ARM processors.

The TimeSeries Datablade, created by Informix in 1998, has been revisited and enhanced, and allowed IBM Informix, jointly with AMT-SYBEX, to literally disintegrate the competition during a very [famous Benchmark on smart electric counters](#). The SQL syntax is extended at each new release, making the syntax gap with other vendors smaller and smaller. More and more admin operations are now executable online.

Multi-Index access method allows the Informix optimizer to choose and use several indexes on one table during a query, which consequently accelerates complex queries where one index only was not very efficient. The STAR JOIN method is a derivate of this features and will give an efficient response to many OLAP queries performance issues.

Informix Warehouse Accelerator: a true innovation!

The most innovative new feature included with 11.70 xC2 is certainly Informix Warehouse Accelerator, also called IBM BLU ACCELERATION. This is a real strategic project of IBM R&D Laboratories where the Informix R&D team has played a major role. Until then, OLAP and Decision Support infrastructure had to reside in separate systems, using huge servers using gigantic disk storage, zillions of CPUs and xillions of memory: extremely expensive infrastructures, also expensive in terms of staff to administrate and permanently create new tables, new indexes or new dataload procedures according to the end users will. Sincerely, the results obtained by spending so much money was really disappointing, queries took hours or even tens of hours before getting a usable result, and any minimal change costed a lot of time and money.

Other solutions, based on appliances, can provide much faster results, but are unfortunately based on expensive to very expensive server technologies and need very highly skilled staff to administrate. Another bad thing is that those systems generally will oblige you to rewrite your DS applications to fit the appliance style. Not every company can / wants to spend 7 figures budgets for their DSS, nor want to be imprisoned by those expensive systems.

Technologies used by IWA are not new ones, but a smart combination of existing ones. The principle of loading the data in RAM(in-memory DB), has been used a long time ago, and it provides stunning performance. A very interesting efficiency level has been observed with the

columnar organization of data(rows become columns at data load time), making the sequential reads much more efficient than they used to be. One of the really nice side effects of this technique is to make indexes unuseful, which will drastically reduce the necessary space for them(no more space needed), but also largely simplify administration tasks because you need no more indexes to create, alter et...

The third technology implemented in IWA is Data Compression. Apart from needing less memory space (according to circumstances from 3 times to more times less), this technique will also reduce consequently the time necessary to read the data:more data in a memory page, all the data will be read faster. This method is generally perfect for DSS queries that will largely apply sequential scans on the data.

I may be reading in your mind your intention to discard the Warehouse Accelerator option because you will have to spend at least a 6 figures amount of money (euros or us \$), and this is where you need to understand a few facts:

- 1) IWA works only on Intel x86 + Linux platforms, which eliminates most of the expensive hardware. With few thousands €/US\$ you can mount a several hundreds of gigabytes datamart. Spend you money in RAM acquisition.
- 2) IWA is very moderate in terms of disk-space, the data is loaded from the Informix OLTP instance either by batch, or in continuous flow.
- 3) IWA has no index, which reduces consequently the necessary RAM space, added to saving staff that used to work a lot of time on accurate indexes creation
- 4) IWA plugs directly to the OLTP Informix instance. This means that, if you allow it, those users can run Accelerated Queries. This also means that absolutely no application change is necessary, nor connection change. Modify nothing and it works... faster!
- 5) How about response times? As a rule of thumb, queries formerly executed in several hours are executed in minutes by IWA. Queries formerly executed in minutes are executed in seconds, or ... second.

Sincerely, even if IWA is not for free, consider that it runs on very affordable servers, and license wise, you can achieve very effective results with a few tens of thousands €/US\$. If you compare with the full cost of ownership of competitors solutions, you will quickly understand that by spending a very moderate amount of money, you will be enabled to radically modify your way of using Decision Support Systems.

Informix 12.10 : simply powerful !

With the rollout of version 12.10 in 2013, IBM confirms a real commitment with the Informix technology, bringing once again major new features. No more smear campaign coming from the IBM executive level, it seems that the product owner has begun understanding its benefits.

An impressive number of new SQL statements is part of the new features, as well as functionality oriented towards auto-administration of the database engine for embedded solutions or just companies that cannot afford a DBA.

Nonetheless, a couple of new features are clearly remarkable. Grid Queries or Sharded Queries consist in querying a table physically located in several of the Flexible Grid servers, either by just adding the GRID keyword in the SQL statement, or by using an environment variable. You do not

have any more to write complex UNION statements, which must eventually be re-written if a new member is added or remove to the Grid.

It is sometimes difficult to believe, if one is used to sensational marketing campaigns from some competitors, that so much technology is so efficiently involved by using a simple environment variable, leading to simpler applications and maintenance cost savings.

The second key feature of version 12.10 is the integration of the unstructured data models called JSON and BSON which is the base of the NoSQL technology. This new feature IS the real open door to the BIG DATA universe, which you will access thru the use of the included 150% compatible MongoDB driver. Informix 12.10 is the unique opportunity to combine all the advantages of the OLTP engine, including Flexible Grid, ACID properties and so many funny things, with the dimension of unstructured data as handle by a NoSQL engine. No need to separate both worlds, you can even execute an SQL join between an SQL 'traditional' table with a NoSQL collection. More savings here: no costly data transfers, no additional DBA to hire with special skills: only IBM Informix can do this as of today!!

IBM Informix can do all this and we were not aware ?

The great characteristic of IBM Informix 12.10 is that is has the capacity to work smoothly on both very small systems (mobile devices/ARM processors) and on complex heterogeneous clustered environments featuring MACH11 and NoSQL features: all of that with one and only one product, no hidden option nor additional cost.

Its robustness, unchanged since the 90's, is the reason why we can find IBM Informix as the core of the IT infrastructure of industries, services or administrations requiring the highest levels of performance, reliability and availability.

Those users appreciate its performance in terms of global cost of ownership, which is materialized by very low staffing and its low system resource consumption leading to a cheaper hardware infrastructure. Using the data compression option can also strongly influence on disk storage costs reduction.

IBM Informix has a discontinued leadership in the IT customer satisfaction enquiries. We can also notice that very often Informix is a requested skillset for DBA job announces. It is to be known that a [recent survey](#) has classified Informix engines in the TOP 10 of the skills most researched by the mobile applications development industry.

Where is Informix? Informix is everywhere, you just don't see it.

Can a company keep using IBM Informix without fearing a sudden discontinuation of the product line by IBM?

IBM has a very clear development roadmap, featuring several major releases, on top of which must be added the 10 years of legal obligation of maintenance after the eventual last release: this is more less putting the date flag by 2024 if IBM suddenly decided tomorrow to discontinue a profitable line of business which is Informix...

So, why would you absolutely want to migrate towards technologies having a higher cost of ownership? Why would you take additional time and cost risks caused by data migration, a new

development coming with its own risks, or even the deployment of a new application about which you do not know whether it will really fit your company's needs?

Have you really thought about a possible rejection of the new system by the end users, because it may be less or not suited and finally less productive than your Informix application that has been bug free for a number of years now?

Are you sure you can assume all those risks, and support your decision with rumors broadcasted by people who just want to take your money, more than having a satisfied customer, achieving end users productivity and working towards your company's profitability?

Today, I can see so many customers that have been loyal to IBM Informix, including distribution worldwide leaders, mission-critical services, government sensitive administrations and just regular companies. Those organizations would not like to be slaves of their IT Infrastructure, but yes, they like their IT infrastructure to serve them.

Many of those organizations have been faithful for more than 25 years. Do you sincerely think this loyalty can only be based on the "Never change a winning team" proverb?